

Requested Patent: WO0079721A2

Title:

SYSTEMS, METHODS AND COMPUTER PROGRAM PRODUCTS FOR
ADJUSTING A TIMEOUT FOR MESSAGE RETRANSMISSION BASED ON
MEASURED ROUND-TRIP COMMUNICATIONS DELAYS ;

Abstracted Patent: WO0079721 ;

Publication Date: 2000-12-28 ;

Inventor(s): SEIFERT JEFFREY D; GROHN OSSII ;

Applicant(s): ERICSSON INC (US) ;

Application Number: WO2000US14265 20000523 ;

Priority Number(s): US19990337182 19990621 ;

IPC Classification: H04L1/00 ;

Equivalents: AU5159400

ABSTRACT:

A round-trip communications delay between two devices that communicate with one another over a communications network is repeatedly measured. An adjustable retransmission timeout is adjusted based upon at least one of the round-trip communications delays that are repeatedly measured. A message that is not acknowledged during the adjustable retransmission timeout is retransmitted. Accordingly, the overall communications speed between the two devices may be increased by accounting for the actual delay that is present in a network at a current time, and basing an adjustable retransmission timeout on this actual delay.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
28 December 2000 (28.12.2000)

PCT

(10) International Publication Number
WO 00/79721 A2

(51) International Patent Classification⁷: H04L 1/00

(21) International Application Number: PCT/US00/14265

(22) International Filing Date: 23 May 2000 (23.05.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/337,182 21 June 1999 (21.06.1999) US

(71) Applicant: ERICSSON, INC. [US/US]; 7001 Development Drive, P.O. Box 13969, Research Triangle Park, NC 27709 (US).

(72) Inventors: GROHN, Ossi, I.; 106 Spindle Creek Court, Apex, NC 27502 (US). SEIFERT, Jeffrey, D.; 3 Kilbreth Avenue, Durham, NC 27713 (US).

(74) Agents: BIGEL, Mitchell, S. et al.; Myers, Bigel, Sibley & Sajovec, P.A., P.O. Box 37428, Raleigh, NC 27627 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AT (utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (utility model), DE, DE (utility model), DK, DK (utility model), DM, DZ, EE, EE (utility model), ES, FI, FI (utility model), GB, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (utility model), SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— Without international search report and to be republished upon receipt of that report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

WO 00/79721 A2

(54) Title: SYSTEMS, METHODS AND COMPUTER PROGRAM PRODUCTS FOR ADJUSTING A TIMEOUT FOR MESSAGE RETRANSMISSION BASED ON MEASURED ROUND-TRIP COMMUNICATIONS DELAYS

(57) Abstract: A round-trip communications delay between two devices that communicate with one another over a communications network is repeatedly measured. An adjustable retransmission timeout is adjusted based upon at least one of the round-trip communications delays that are repeatedly measured. A message that is not acknowledged during the adjustable retransmission timeout is retransmitted. Accordingly, the overall communications speed between the two devices may be increased by accounting for the actual delay that is present in a network at a current time, and basing an adjustable retransmission timeout on this actual delay.

**SYSTEMS, METHODS AND COMPUTER PROGRAM PRODUCTS
FOR ADJUSTING A TIMEOUT FOR MESSAGE RETRANSMISSION
BASED ON MEASURED ROUND-TRIP COMMUNICATIONS DELAYS**

Field of the Invention

This invention relates to communications systems, methods and computer program products, and more particularly to systems, methods and computer program products for communicating between two devices that are connected by a
5 communications network.

Background of the Invention

Communications networks are widely used to communicate among devices that are attached thereto. As is well known to those having skill in the art, the devices
10 may include mainframe computers, midrange computers, personal computers, workstations, terminals and/or telecommunications devices, such as cellular radiotelephones, satellite radiotelephones, personal digital assistants and/or other wireless communications devices. The network may include a Local Area Network, a Wide Area Network, the Internet, an intranet, a cellular/satellite radiotelephone
15 communications network and/or other wired or wireless communications networks. Communications may take place in a peer-to-peer, client-server or other environments. Figure 1 is a simplified block diagram of conventional network communications between Device 1 100 and Device 2 120 over a communications network 110.

20 When two devices 100 and 120 communicate between one another over a communications network 110, the messages between the devices may be lost due to errors in the network and/or other reasons. When messages are lost, they generally are retransmitted. In order to determine whether a message has been lost, a receiving device generally transmits an acknowledgement message to a transmitting device for
25 each received message. If the transmitting device does not receive the acknowledgement message after a fixed time elapses, it can retransmit the original message.

The fixed time, also referred to as a retransmission timeout, generally is set to be long enough to account for all possible delays that may be present in the network. By accounting for all possible delays, the number of messages that are needlessly retransmitted can be reduced and preferably can be eliminated. Unfortunately, when
5 an error occurs in the network and a message sent by a first device is not received by a second device, the long retransmission timeout may produce long communications delays between the two devices. These delays can degrade communications between the devices.

10 Summary of the Invention

It is therefore an object of the present invention to provide improved systems, methods and computer program products for communicating between two devices that are connected by a communications network.

It is another object of the present invention to provide systems, methods and
15 computer program products for communicating between two devices that are connected by a communications network, which can reduce the retransmission timeout that controls retransmission of a message that is not received.

These and other objects are provided according to the present invention by systems, methods and computer program products for communicating between two
20 devices that are connected by a communications network, by repeatedly measuring a round-trip communications delay between the two devices over the communications network, and adjusting an adjustable retransmission timeout based upon at least one of the round-trip communications delays that are repeatedly measured. A message that is not acknowledged during the adjustable retransmission timeout is
25 retransmitted. The adjustable retransmission timeout preferably is based upon the last measured round-trip communications delay. Accordingly, the overall communications speed between the two devices may be increased by accounting for the actual delay that is present in the network at a given time, and basing an adjustable retransmission timeout on this actual delay. Improved systems, methods and
30 computer program products for communicating between two devices that are connected by a communications network thereby can be provided.

In particular, according to the present invention, a first message is transmitted (sent) at a first (transmit) time and an acknowledgement of the first message is received at a second (receive) time. A time difference between the second time and

the first time is then determined, to thereby obtain the adjustable retransmission timeout. Transmission, reception and determination of a time difference are repeatedly performed, preferably on a periodic basis, to thereby measure the actual round-trip delay in the network, and thereby obtain a suitable value for the timeout.

- 5 The first message may be a unique delay measurement message that is used to measure delay and that does not contain communications data. The delay measurement message may contain a timestamp therein that indicates the first time. Alternatively, the transmitting device may maintain a record of the time that the first message was transmitted. The delay measurement message also preferably contains
- 10 padding therein, so that the delay measurement is at least as long as a longest communication message between the two devices. Thus, although the round-trip communications delay is adjusted based on the actual performance of the network, the performance can represent the actual worst case performance for a longest message.

- The acknowledgement message may be an echo back of the first message.
- 15 Alternatively, the acknowledgement message may be a unique acknowledgement message.

- In yet another alternative, the first message is not a unique delay measurement message. Rather, the first message may be a preexisting communications message that is sent from the first device to the second device. The communications message
- 20 may include a timestamp, or the time that the first message was sent may be recorded in the first device.

- The adjustable time may correspond identically to the last measured round-trip delay. Alternatively, the adjustable time may be based on an average of a predetermined number of last round-trip delay measurements. In yet another
- 25 alternative, a guard band may be added to the last measured round-trip delay to account for potential changes since the last measured delay time. Other techniques for determining the adjustable time may be provided.

- Accordingly, the overall communications speed between the devices may be increased since the retransmission may be performed as soon as possible. A higher
- 30 number of errors also may be able to be tolerated by the communications network, since the number of retransmissions during a given period of time may be increased. By automatically adapting to changed communication network delays, improved systems, methods and computer program products for communicating between two devices that are connected by a communications network may be provided.

Brief Description of the Drawings

Figure 1 is a block diagram of conventional communications between devices over a network.

5 Figure 2 is a block diagram of a device according to an embodiment of the present invention.

Figure 3 is a flowchart illustrating communications between devices according to an embodiment of the present invention.

10 Figure 4 is a timing diagram illustrating communications between devices according to an embodiment of the present invention.

Figure 5 is a flowchart illustrating detailed operations for repeatedly measuring a round-trip delay according to an embodiment of the present invention.

Figure 6 is a flowchart illustrating operations for retransmitting a message according to an embodiment of the present invention.

15

Detailed Description of Preferred Embodiments

The present invention now will be described more fully hereinafter with reference to the accompanying drawings, in which preferred embodiments of the invention are shown. This invention may, however, be embodied in many different
20 forms and should not be construed as limited to the embodiments set forth herein; rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art. Like numbers refer to like elements throughout.

Various aspects of the present invention are illustrated in detail in Figures 2-6,
25 including flowchart illustrations. It will be understood that individual blocks of the figures, and combinations of blocks in the figures, can be implemented by computer program instructions. These computer program instructions may be provided to a processor or other programmable data processing apparatus to produce a machine, such that the instructions which execute on the processor or other programmable data
30 processing apparatus create means for implementing the functions specified in the block or blocks. These computer program instructions may also be stored in a computer-readable memory that can direct a processor or other programmable data processing apparatus to function in a particular manner, such that the instructions

stored in the computer-readable memory produce an article of manufacture including instruction means which implement the functions specified in the block or blocks.

Accordingly, blocks of the figures support combinations of means for performing the specified functions, combinations of steps for performing the specified functions and program instruction means for performing the specified functions. It will also be understood that individual blocks of the figures, and combinations of blocks in the figures, can be implemented by special purpose hardware-based computer systems which perform the specified functions or steps, or by combinations of special purpose hardware and computer instructions.

Referring now to Figure 2, a block diagram of a device 100' according to an embodiment of the present invention will now be described. The device 100' may communicate with a second device 120 over a network 110 in a manner that was illustrated in Figure 1. As shown in Figure 2, the device 100' includes dynamic communications delay measurement section 210 that repeatedly measures a round-trip communications delay between the two devices over the communications network. A transmission section 220 transmits a message from the device 100' to the second device 120 over the network 110. A retransmission section 230 retransmits a message that is not acknowledged during an adjustable retransmission timeout after the initial transmission by the transmission section 220. The adjustable retransmission timeout is determined by a retransmission timeout adjustment section 200, based upon at least one of the repeated round-trip communications delays that are repeatedly measured by the dynamic communications delay measurement section 210.

Figure 3 illustrates an overview of communicating between devices according to the present invention. As shown at Block 310, round-trip communications delays between the two devices over the communications network are repeatedly measured. For example, as shown in Figure 4, at Time 1, a first message in the form of a unique Delay Measurement Message is transmitted from the transmitting device to the receiving device. The Delay Measurement Message includes a timestamp (Time 1) therein. The receipt of the Delay Measurement Message is acknowledged by the receiving device echoing back the Delay Measurement Message. The echoed back Delay Measurement Message arrives at the transmitting device at a second time Time 2. A time difference between the second time and the first time is then determined in order to determine a retransmission timeout.

It will be understood by those having skill in the art that a unique measurement message need not be transmitted in order to measure the round-trip communications delay. Rather, a conventional data communications message may be sent and the response to the data communications message may be measured. When a
5 conventional data communications message is transmitted, the response may be a different response message rather than an echo back. It also will be understood that the Delay Measurement Message or the conventional message need not include a timestamp therein. Rather, the time that the message was transmitted may be recorded at the transmitting device and need not be included in the message itself.

10 Referring back to Figure 3, at Block 320, the adjustable retransmission timeout is adjusted based upon at least one of the round-trip communications delays that are repeatedly measured. Preferably, the retransmission timeout is adjusted based upon the last one of the round-trip communications delays that are repeatedly measured, in order to provide a timeout that reflects a current state of communications
15 between the devices over the network. It will be understood that although the adjustable retransmission timeout preferably is adjusted in the device 100', it may be adjusted outside the device 100'.

Still referring to Figure 3, at Block 330, a message that is not acknowledged during the adjustable timeout that is based upon the measured round-trip delay is
20 retransmitted. In detail, referring again to Figure 4, a Communications Message is transmitted by the transmission section 220 of Figure 2. If an acknowledgement message is not received by the transmitting device within the adjustable retransmission timeout that is based on at least one of the round-trip communications delays that are repeatedly measured, then the Communications Message is
25 retransmitted from the transmitting device to the receiving device.

Accordingly, a device measures a round-trip delay by repeatedly and preferably periodically transmitting a message with a current timestamp to another device that responds by echoing the measurement signal back to the transmitting device. The transmitting device then determines the round-trip delay by subtracting
30 the received timestamp value from the current time. This retransmission timeout value is an indicator of how long the device should wait for an acknowledgement from the other device. If a device does not receive the acknowledgement before the end of the retransmission timeout, it retransmits the original message.

Repeated round-trip delay measurements are performed because the round-trip delay over the communications network can change over time due to altered routing, traffic changes, changed transmission equipment, temperature variations and/or other reasons. Moreover, if new functions are added to a device, the response time may change. The repeated delay measurement allows the device to optimize the retransmission timeout value automatically if the total round-trip delay changes.

Referring now to Figure 5, detailed operations for repeatedly measuring the round-trip delay (Block 310 of Figure 3) now will be described. At Block 510, the round-trip communications delay, denoted herein by *LinkDelay*, is initialized to a maximum allowable delay, denoted herein by *MaxDelay*, based on the type of connection between the two devices. At Block 515, a delay measurement period, denoted herein by *DelayPeriod*, is selected. The delay measurement period defines the periodicity at which the round-trip communications delay between the two devices over the communications network is measured. The *DelayPeriod* may be selected based upon the type of connection, such as local area network, wide area network or the Internet, and can vary from seconds or less to days or more.

Continuing with the description of Figure 5, at Block 520, operations wait until *DelayPeriod* expires or connection between the devices is established or reestablished. At that time, as shown at Block 525, the timestamp, denoted herein by *TimeStamp*, is read, for example from a free-running counter that is incremented every *r* milliseconds (ms). Then, at Block 530, the Delay Measurement Message is transmitted to the receiving device. The Delay Measurement Message may have the following format:

Message Type	Data		Checksum
<i>DelMeas</i>	<i>TimeStamp</i>	<i>Padding</i>	<i>CS_DelMeas</i>

As shown, the Delay Measurement Message may have a *DelMeas* identifier that identifies the message. The Delay Measurement Message also may include the *TimeStamp* and a checksum. A *Padding* field may also be included. The length of the *Padding* field preferably is selected so that the delay measurement message is the longest message that is sent between the devices. In this manner, the maximum round-trip delay may be measured.

Continuing with the description of Figure 5, at Block 535, the receiving device recognizes the Delay Measurement Message based on the message type (*DelMeas*) and echoes it back. A test is made as to whether the echo back is received within the maximum allowable delay period *MaxDelay*. If not, a new message is sent beginning
 5 at Block 525. If yes, a calculation is made as to whether the calculated checksum is equal to the received checksum at Block 540. If not, a new message is sent beginning at Block 525.

If the echo back was received within the maximum allowable delay period and the checksum is valid, then at Block 545, the current counter value, denoted by
 10 *CurrentTime*, is read. The round-trip delay (*LinkDelay*) is then set at Block 550 using the following equation:

$$LinkDelay = CurrentTime - TimeStamp.$$

The adjustable retransmission timeout may be set to this last measured round-trip communications delay. Alternatively, the adjustable time may be based on an
 15 average of a predetermined number of last round-trip delay measurements. In yet another alternative, a guard band may be added to the last measured round-trip delay to account for potential changes since the last measured delay. Other techniques for determining the adjustable retransmission timeout may be provided.

It will be understood that the size (n bits) and the time resolution (r ms) of the
 20 counter may determine the minimum and maximum measurable delays as follows:

$$\begin{aligned} DelayMin &= r \text{ ms} \\ DelayMax &= (2^n - 1) \times r \text{ ms} \end{aligned}$$

Accordingly, if $n = 16$ and $r = 1$ ms, then $DelayMax = (2^{16} - 1) \times 1 \text{ ms} = 65.535 \text{ s}$. It also will be understood that if link delay is less than 0, this generally means that the
 25 counter has rolled over, so that link delay is set to $2^n + LinkDelay$.

Referring now to Figure 6, detailed operations for retransmitting a message that is not acknowledged during an adjustable retransmission time based upon the measured round-trip delay (Block 330 of Figure 3) now will be described. As shown in Block 605, a retransmission counter, denoted by *RetrCntr*, is initialized to the
 30 maximum number of allowable retransmissions denoted by *MaxRetr*. The value of *MaxRetr* may depend on the application. Then at Block 610, a timestamp denoted *TimeStamp* is read from the free-running counter. At Block 615, a message is transmitted to the receiving device. The message may have the following form, although other forms may be used:

Message Type	Data	Checksum
<i>Sig_X</i>	<i>Data_X</i>	<i>CS_SigX</i>

If the receiving device receives the message and if the calculated checksum is equal to the received checksum (*CS_SigX*), it sends back an acknowledge message
 5 that may have the following format:

Signal Type	Checksum
<i>Ack_X</i>	<i>CS_AckX</i>

Other acknowledgement messages may be used and/or a reply message may be sent.

Referring now to Block 620, a *CurrentDelay* is calculated by reading the
 10 current counter value, referred to as *CurrentTime*. The *CurrentDelay* then is calculated as:

$$CurrentDelay = CurrentTime - TimeStamp.$$

It also will be understood that if the value of *CurrentDelay* is less than 0, this generally means that the counter has rolled over and *CurrentDelay* is calculated by
 15 the following relationship:

$$CurrentDelay = 2^n + CurrentDelay.$$

Then, at Block 625, a test is made as to whether the *CurrentDelay* is less than the retransmission timeout that was set as a result of the operations of Figure 5. If the *CurrentDelay* is not less than the retransmission timeout that was set, then at Block
 20 645, a determination is made as to whether the retransmission counter is greater than zero. If the retransmission counter is greater than zero, then the retransmission counter is decremented at Block 650 and the message is retransmitted beginning at Block 610. On the other hand, if the retransmission counter is zero at Block 645, there are no retransmission attempts left. Error handling is then performed beginning
 25 at Block 655, using conventional error handling procedures.

Returning again to Block 625, if the *CurrentDelay* is less than the retransmission timeout, then a test is made at Block 630 as to whether an acknowledgement has been received. If the acknowledgement has not been received, the *CurrentDelay* is recalculated until the acknowledgement is received or the
 30 *CurrentDelay* is longer than the delay period at Blocks 625 and 630. If the

acknowledgment is received at Block 630 prior to expiration of the retransmission timeout, then at Block 635, a test is made as to whether the calculated checksum is equal to the received checksum *CS_AckX*. If the checksums are equal, the acknowledgment message was received in time and is valid. The message was
5 successfully transmitted as indicated at Block 640 and a new message is transmitted by again performing the operations beginning at Block 605. On the other hand, if the checksum was not valid, then the operations beginning at Block 645 are performed again to determine whether the message can be retransmitted or if there are no more retransmission attempts left.

10 Accordingly, the retransmission timeout may be optimized based on the round-trip delay time of the communications network approximately contemporaneous with the transmission of a message. Realistic timeouts therefore may be used, that are neither too long nor too short, based on the current network performance.

15 In the drawings and specification, there have been disclosed typical preferred embodiments of the invention and, although specific terms are employed, they are used in a generic and descriptive sense only and not for purposes of limitation, the scope of the invention being set forth in the following claims.

What is Claimed is:

1. A method of communicating between two devices that are connected by a communications network, the method comprising the steps of:
repeatedly measuring a round-trip communications delay between the two devices over the communications network;
5 adjusting an adjustable retransmission timeout based upon at least one of the round-trip communications delays that are repeatedly measured; and
retransmitting a message that is not acknowledged by the second device during the adjustable retransmission timeout.
2. A method according to Claim 1 wherein the repeatedly measuring step comprises the following steps that are repeatedly performed:
transmitting a first message at a transmit time;
receiving an acknowledgement of the first message at a receive time; and
5 determining a time difference between the receive time and the transmit time to thereby obtain the adjustable retransmission timeout.
3. A method according to Claim 2:
wherein the first message is a delay measurement message that is used to measure delay and that does not contain communications data; and
wherein the acknowledgement is an echo back of the first message.
4. A method according to Claim 3 wherein the delay measurement message contains a timestamp therein that indicates the first time.
5. A method according to Claim 4 wherein the delay measurement message also contains padding therein so that the delay measurement message is at least as long as a longest communication message between the two devices.
6. A method according to Claim 2 wherein the first message is a data communication message between the two devices that also is used to determine the time difference.

7. A method according to Claim 1 wherein the adjusting step comprises the step of adjusting an adjustable retransmission timeout based upon a last one of the round-trip communications delays that are repeatedly measured.

8. A method according to Claim 1 wherein the retransmitting step comprises the steps of:

awaiting receipt of an acknowledgement of the first message during the adjustable retransmission timeout; and

5 retransmitting the first message if the acknowledgment is not received during the adjustable retransmission timeout.

9. A system for communicating between two devices that are connected by a communications network, the system comprising:

means for repeatedly measuring a round-trip communications delay between the two devices over the communications network;

5 means for adjusting an adjustable retransmission timeout based upon at least one of the round-trip communications delays that are repeatedly measured; and

means for retransmitting a message that is not acknowledged during the adjustable retransmission timeout.

10. A system according to Claim 9 wherein the means for repeatedly measuring comprises:

means for transmitting a first message at a transmit time;

5 means for receiving an acknowledgement of the first message at a receive time; and

means for determining a time difference between the receive time and the transmit time to thereby obtain the adjustable retransmission timeout.

11. A system according to Claim 10:

wherein the first message is a delay measurement message that is used to measure delay and that does not contain communications data; and

wherein the acknowledgement is an echo back of the first message.

12. A system according to Claim 11 wherein the delay measurement message contains a timestamp therein that indicates the transmit time.

13. A system according to Claim 12 wherein the delay measurement message also contains padding therein so that the delay measurement message is at least as long as a longest communication message between the two devices.

14. A system according to Claim 10 wherein the first message is a data communications message between the two devices that also is used to determine the time difference.

15. A system according to Claim 9 wherein the means for adjusting comprises means for adjusting an adjustable retransmission timeout based upon a last one of the round-trip communications delays that are repeatedly measured.

16. A system according to Claim 9 wherein the means for retransmitting comprises:

means for awaiting receipt of an acknowledgement of the first message during the adjustable retransmission timeout; and

5 means for retransmitting the first message if the acknowledgment is not received during the adjustable retransmission timeout.

17. A computer program product for communicating between two devices that are connected by a communications network, the computer program product comprising a computer-readable storage medium having computer-readable program code means embodied in the medium, the computer-readable program code means
5 comprising:

computer-readable program code means for repeatedly measuring a round-trip communications delay between the two devices over the communications network;

computer-readable program code means for adjusting an adjustable retransmission timeout based upon at least one of the round-trip communications
10 delays that are repeatedly measured; and

computer-readable program code means for retransmitting a message that is not acknowledged during the adjustable retransmission timeout.

18. A computer program product according to Claim 17 wherein the computer-readable program code means for repeatedly measuring comprises:

computer-readable program code means for transmitting a first message at a transmit time;

5 computer-readable program code means for receiving an acknowledgement of the first message at a receive time; and

computer-readable program code means for determining a time difference between the receive time and the transmit time to thereby obtain the adjustable retransmission timeout.

19. A computer program product according to Claim 18:

wherein the first message is a delay measurement message that is used to measure delay and that does not contain communications data; and

wherein the acknowledgement is an echo back of the first message.

20. A computer program product according to Claim 19 wherein the delay measurement message contains a timestamp therein that indicates the transmit time.

21. A computer program product according to Claim 20 wherein the delay measurement message also contains padding therein so that the delay measurement message is at least as long as a longest communication message between the two devices.

22. A computer program produce according to Claim 18 wherein the first message is a data communications message between the two devices that also is used to determine the time difference.

23. A computer program product according to Claim 17 wherein the computer-readable program code means for adjusting comprises computer-readable program code means for adjusting an adjustable retransmission timeout based upon a last one of the round-trip communications delays that are repeatedly measured.

24. A computer program product according to Claim 17 wherein the computer-readable program code means for retransmitting comprises:

computer-readable program code means for awaiting receipt of an acknowledgement of the first message during the adjustable retransmission timeout;

5 and

computer-readable program code means for retransmitting the first message if the acknowledgment is not received during the adjustable retransmission timeout.

25. A device that communicates with a second device over a communications network, the device comprising:

a dynamic communications delay measurement system that repeatedly measures a round-trip communications delay with the second device over the
5 communications network;

a retransmission timeout adjuster that adjusts an adjustable retransmission timeout based upon at least one of the round-trip communications delays that are repeatedly measured; and

a message retransmitter that retransmits a message that is not acknowledged
10 by the second device during the adjustable retransmission timeout.

26. A device according to Claim 25 wherein the dynamic communications delay measurement system comprises:

a transmitter that transmits a first message at a transmit time;

a receiver that receives an acknowledgement of the first message at a receive
5 time; and

a calculator that determines a time difference between the receive time and the transmit time to thereby obtain the adjustable retransmission timeout.

27. A device according to Claim 26:

wherein the first message is a delay measurement message that is used to measure delay and that does not contain communications data; and

wherein the acknowledgement is an echo back of the first message.

28. A device according to Claim 27 wherein the delay measurement message contains a timestamp therein that indicates the first time.

29. A device according to Claim 28 wherein the delay measurement message also contains padding therein so that the delay measurement message is at least as long as a longest communication message between the device and the second device.

30. A device according to Claim 26 wherein the first message is a data communications message between the device and the second device that also is used to determine the time difference.

31. A device according to Claim 25 wherein the retransmission timeout adjuster adjusts the adjustable retransmission timeout based upon a last one of the round-trip communications delays that are repeatedly measured.

32. A device according to Claim 25 wherein the message retransmitter awaits receipt of an acknowledgement of the first message during the adjustable retransmission timeout, and retransmits the first message if the acknowledgment is not received during the adjustable retransmission timeout.

FIG. 1
PRIOR ART




FIG. 2

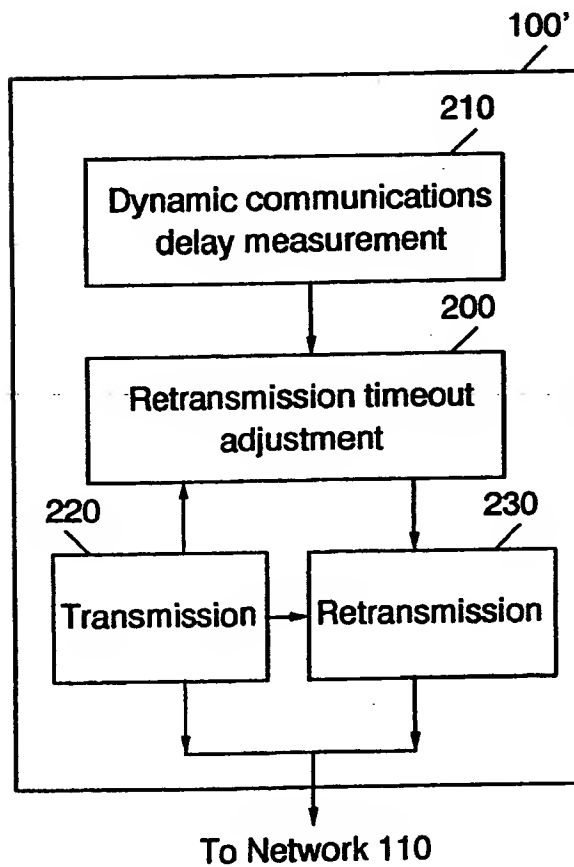


FIG. 3

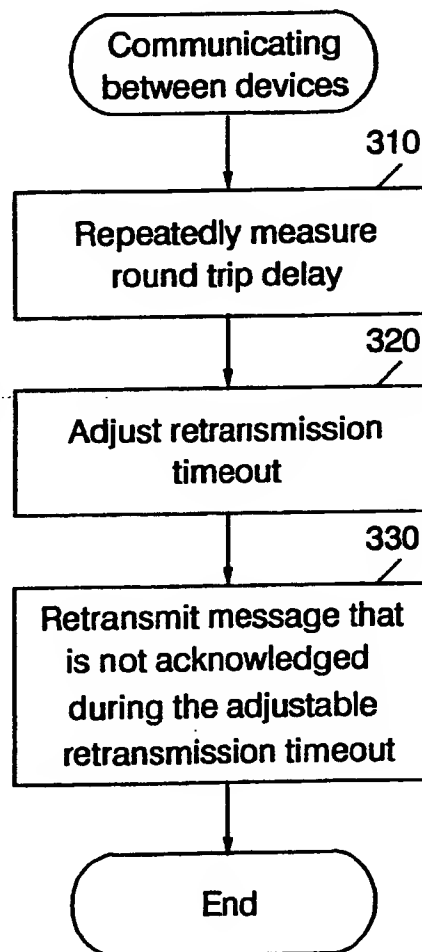
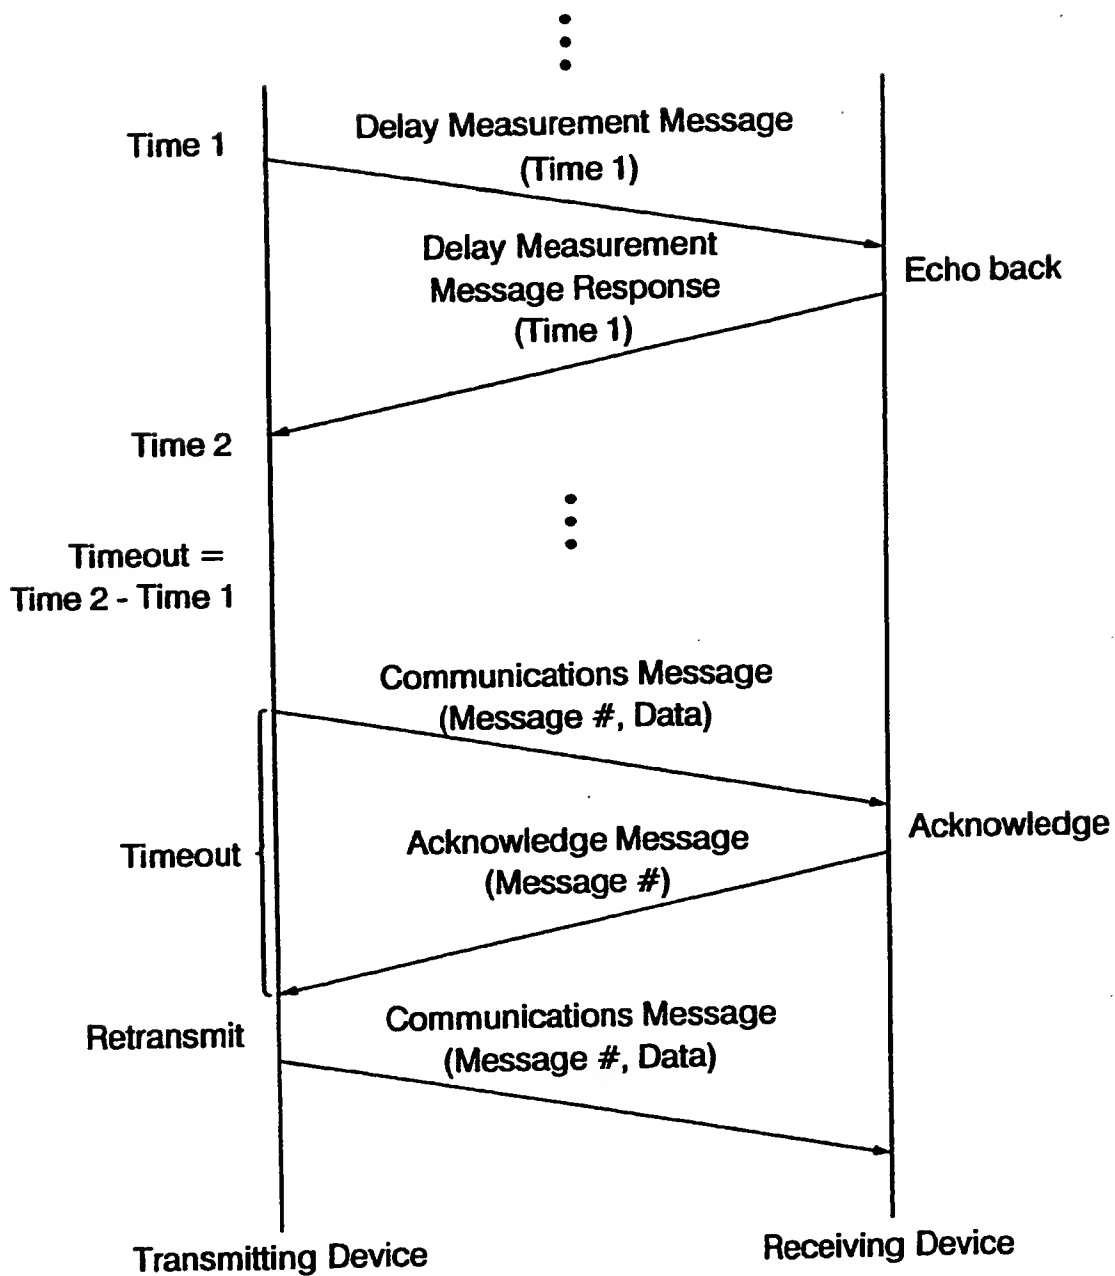


FIG. 4






FIG. 5

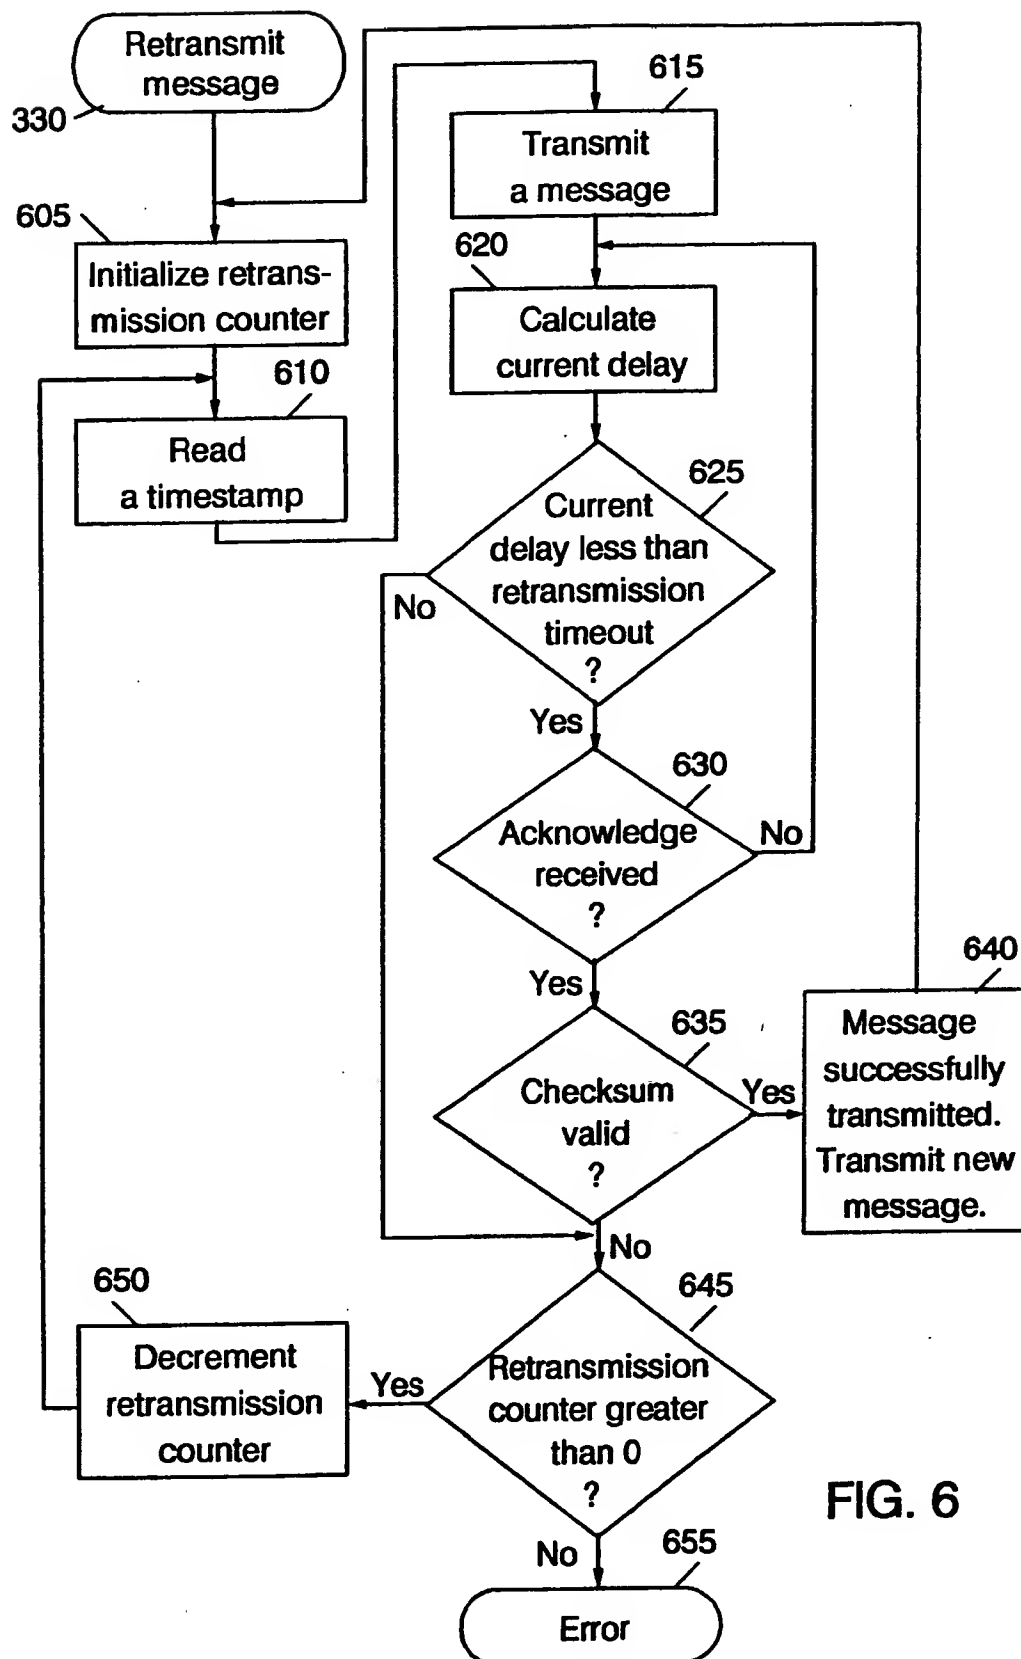


FIG. 6